The Impact of Internet Services on the Telecommunications Infrastructure

Russ Bell, P.E.
Director, Communications Technology
Advanced Micro Devices

AMD is on the leading edge of the collision between computation and communications. Although we are well-known for our position as the number two supplier of 32 bit microprocessors, we are less known for our market ranking in communications as over 40 percent of our sales revenues are derived from the communications market. AMD is the leading non-captive supplier of silicon for central office line cards, having shipped over 100 million, DSP based solutions to 70 countries around the world that provide telephone connectivity and internet access for millions. AMD is also the leading ethernet circuits supplier providing both hub and NIC products for LAN connectivity and ***intranet*** access. AMD blazed the ISDN trail in 1984 with a complete series of basic rate devices for local loop and customer premises equipment that provided unparalleled performance in the industry. Today, our CPE device remains one of the premier performers in the industry. AMD also led the development of high speed LAN technology by developing the most advanced FDDI chip set on the market and by our leadership in the ANTC (advanced networking Test Center) which collapsed the time for vendors to get products to market and assured interoperability between competing solutions. AMD is currently marketing advanced wireless solutions in the 802.11 (wireless LAN) and 900 MHz. cordless ISM bands. So, as you can see, AMD is more than a microprocessor company; we are a computation and communications company. This unique perspective of the industries allows us to take a more balanced view of emerging markets and allows us to more carefully target investments that will enable the next generation of computing for the masses, the fifth generation of computing that merges computing with the vast resources of the public communications network. To that end, it is important to understand more than the requirements of our customers. AMD must understand the requirements of our customer's customer. Specifically, we must understand the business models, technology issues, and economics of the service providers, RBOCs, MSOs, PCS providers, etc. Only then will we be able to craft the optimal solutions for the industry that will deliver the next generation of computing to the mass market.

The Internet

The internet has proven to be one of the most amazing stories of the decade and will be one of the keys to the fifth generation of computing. It is hard to conceive how a network developed by DARPA for interconnecting research institutions around the world could evolve into the "stock-bolstering," headline grabbing machine that we have witnessed over the past two years. It seems that everyone has jumped on the bandwagon, from banks to bagels, you can find it on the net. The introduction of a reasonable web GUI (graphical user interface - MOSAIC) marked the beginning of the internet gold rush. Novice users were now able to navigate the web with little more than a few mouse clicks. ISPs (internet service providers) began springing up around the world and users began going 'on-line' at astronomical rates. The rest of the story is history.

There are, however, some snags that formed along the way. Specifically, the entire internet infrastructure is formed by a patchwork of legacy equipment, systems, and protocols. As in many defacto-driven scenarios, it is doubtful that one would consciously design a modern internet infrastructure with the systems that are in place today. Nevertheless, the world is "stuck" with billions of dollars of investments in infrastructure and billions of lines of code; this will not change overnight. Initially, the rapid internet growth rates were somewhat problematic. Today, the problems are becoming very serious.

What is this Infrastructure Stuff Anyway?

In the beginning, internet growth was fueled by corporate and institutional subscribers. These surfers were connected to the internet via higher speed, dedicated transmission facilities (T1s, etc.) to their ISP. As the success of the internet caught on, the user base migrated from the office to the home and small office environment (SOHO). The predominant communications link to and from these SOHO customers is the telephone line. This line connects to the phone company's central office switching facility which is essentially a huge computer that connects callers to other callers (or ISPs). The telephone system is arguably the largest machine in the world. There are over 700 million telephone lines distributed around the world with an estimated undepreciated value of over 600 billion dollars. Like any great engineering effort, the telephone system is optimized for the task at hand. Specifically, the telephone switching centers are designed to accommodate 4 minute phone calls that occur at a peak rate of 4 calls per hour using 4 kHz. of bandwidth. It is this statistical tightrope that engineers walked in the design phase of the world wide telecommunications system. Unfortunately, the internet has broken the machine, at least temporarily.

The balance of the infrastructure (which consists of trunks from the CO to an ISP) provides connection to the internet backbone using a high speed link (typically T1s, occasionally a DS3) via one of the many NAPs (network access points) around the US. At some point, your internet surfing request arrives at a host computer and the transfer of information begins. As in all communications systems, protocols are used to transfer the information across the network; the internet standard is TCP IP. Those famous "HTTP" protocols run on top of TCP IP. Then there is the graphical nature of web pages that are retrieved over the network. Most of these are HTML (Hyper text Mark-up Language) documents. HTML provides a way to represent information on the web (WWW) so that it can be reconstructed at the requester's PC. Ironically, all of these pieces of the puzzle are contributing to the poor performance of the web. However, it is the very presence of these 'tools' that has propelled the internet into the headlines, for without these loose knit tools, there would be no internet as we know it today.

The First Link

The first signs of internet meltdown began to occur in the local telephone system. Not unlike a PC with insufficient memory, the CO began to run out of capacity. This is caused by WEB browsers who tend to surf for long periods of time. In fact, there are estimates that the average internet 'call' may last as long as 30-40 minutes. It doesn't take many surfers to quickly trash the well-honed statistical models of the telephone companies which expects 3-4 minute phone calls with 3-4 calls per hour. The result is that annoying '***fast busy***' signal that you may have encountered when you attempt to log onto the network. The CO switch is simply out of capacity. This phenomenon is very similar to the traffic patterns that exist during Mothers Day, hurricanes, earthquakes, etc.

Unfortunately, upgrading the switch capacity is MUCH harder than adding additional memory to your PC. The technical term that is often used to describe this problem is ***class five switch blocking***. These problems can certainly be fixed; all it takes is money. The growth of internet traffic from consumers to ISPs through the central office has created a traffic problem for the telephone company while generating revenue for the ISP. The calls to the ISP are typically local calls; this translates to ***" free for the consumer" (note that this is not typically true in Europe as they tend to bill by the minute).*** It is hard to envision a scenario in which the RBOCs will invest in network improvements that provide little or no profitability incentives for them, but improve the access to ISPs (who are potential competitors). Therefore, the problem is unlikely to be resolved easily.

Why is the net so slow?

Many call it the World-Wide-Wait (WWW). The number one internet complaint in a recent Georgia Tech internet survey is ***speed of response.*** Ergonomic studies indicate that a sub 3 second response time is required to maintain a user's attention. Delays of greater than 10 seconds result in serious 'brain fade.' Internet response time is clearly the "hot button" of almost all web surfers and one that must be resolved from a systematic approach.

Viewing a Web Page

The simple act of viewing a web page creates the majority of traffic on the internet today.
Typically, the viewing of a home page requires approximately 12 packets being transferred over 5 separate networks. The sequence is:

1) The LOCAL computer sends information over local MODEM to the CO.

2) The CO transfers the call to the ISP, typically via a T1 trunk.

3) The ISP connects to a national backbone provider (MCI, Sprint, BBN, etc.).

4) The backbone provider passes your packets through one of the peering locations which connects your message to another backbone provider.

5) The second backbone provider connects with the web site's ISP.

6) The ISP connects to the appropriate server.

7) This entire process is repeated in the reverse direction

Many believe that "pumping up the data rate from 28.8 kbps to ISDN (128 kbps) or ADSL (6 Mbps / 640 kbps) rates will "turbo the net," thus improving the overall response time of internet access. Much has been written about the cable MODEM and the "lightning fast response times" that 10 Mbps cable technology will bring to the internet surfer. Unfortunately, as is often the case, systemic problems require systemic solutions. **The entire system delay must be addressed if we are to see truly fast internet access.**

Hurry up and Wait

There are many components associated with web delay including page complexity, distance that packets must travel, the number of router hops, the power of the server at the web site, the protocol settings and configuration and access line rates. As figure 14 indicates, there are several components of delay; addressing a single portion of the problem will not yield optimal results.

Figure 7 also highlights the relationship between access rates and home page delivery times. The vertical bands represent various transmission rates ranging from 28.8 kbps to 10 Mbps cable modem technology. It should be obvious from this graph that the replacement of your 28.8 kbps modem with a 10 Mbps cable modem will do little to enhance your web browsing experience. However, this is not to say that high speed data services are not needed. There are a variety of non-internet applications that require high speed data services including videoconferencing, narrowcasting, entertainment, etc. Furthermore, once the entire internet infrastructure has been upgraded, the entire surfing experience will improve radically.

This graph is based on the delivery of 50 k bytes of text and graphical information depicting a typical home page. The solid lines depict ‘typical’ web systems; the dotted lines represent the same 50 k bytes of information delivered by ‘more efficient’ web pages (i.e. less overhead). Then, in each of these two scenarios, the graph provides additional detail on standard internet, upgraded internet, and desktop caching.

An Upgraded internet (for purposes of this discussion) consists of a corporate server that is replicated in several regional areas. Server replication on a geographical basis provides substantial performance gains for surfers as the location of the information is physically “closer” to the requester, thus minimizing the number of router “hops” as well as potentially minimizing the number of peering requests that must be made when crossing from one IXC’s backbone to another. The connection to the ISP is assumed to be T3 or greater and protocols are optimized for throughput.

A desktop cache system configures the navigator to cache locally. These curves also assume standard internet access (up to 10 Mbytes of local storage may be required). Overall, caching is the one biggest improvement that can be made to improve overall web performance.

Ironically, the much touted network computer (the diskless, low cost desktop PC replacement for internet access) will be “inter and intranet challenged” if there is not a reasonably sized local cache suitable for web page use.

Protocol Issues - Hypertext and more

The HTTP command that many of us are familiar with is the abbreviation for hypertext Transfer Protocol. HTTP is transferred on top of TCP IP and TCP packets are routed across the network to the final destination. Unfortunately, HTTP is not a very efficient transfer protocol. A requesting client will receive an HTTP document via an HTTP request. Once the document is received, the client may discover that there are several links inside the HTML document. A typical document may consist of 6.4 k bytes plus seven images per page. Each of these links may require a separate TCP-IP connection to be established between the client and the server, and then the fun really begins.

At the Server

Things are not much better at the server side of the network. The server is required to maintain state information for each TCP connection that is 'open.' When a connection in TCP is closed, a handshaking procedure requires that the remote end of the link acknowledge the closure. It is possible that the acknowledgment could be lost and the local user would time out and retransmit the final ***close packet*** to the remote location. The TCP specification requires that the link remain in the TIME_WAIT state for 2MSL (maximum segment lifetime) after the local user has transmitted the active close segment to the remote end. TCP specifies this time as 4 minutes. The secondary affect of this TIME_WAIT state is on the OS. The OS can not release the socket associated with this connection until the timer has expired; this includes any resources associated with the OS including TCBs (transmission control blocks) which save unacknowledged packets, control information, etc., about the current connection. This precaution is required to prevent any sort of data corruption caused by sequence number re-use in TCP. The sequence numbers are used by the TCP protocol to re-assemble various packets as they arrive to the termination point. Occasionally, these packets may be lost and the receiving entity can determine which packets are lost by examining consecutive sequence numbers; the receiver can simply request a re-transmission of the lost packet. There are some transmission situations that result in **delayed duplicates**, whereby congestion on a route could cause a packet to be delayed for some time before it ultimately arrives at he destination address; in the interim time (while the packet was lost), the receiver detected the loss and requested a re-transmission. The retransmission was received and the TCP connection was closed. No segments that arrive during the TIME_WAIT state will be accepted. These would be the segments that were the "**delayed duplicates".** Once the timer has expired, it will not be possible for any other existing duplicates to be on the network. Any subsequent connection that is established using the now-released socket will not suffer from receiving segments associated with another connection.

Furthermore, the TCB can not be released for several minutes after the actual TCP IP connection has been terminated. This may result in an exhaustion of resources at the server end as the TCBs are maintained for several minutes even though they are not being utilized. Perhaps you are familiar with the **"host contacted - awaiting response**" message that is typically caused by TCB exhaustion. The client may have received the web page and all associated text and graphical data, but the server may not release the TCBs back to the OS for several minutes (depending on the value of the MSL), effectively blocking the use of those resources by other surfers. The typical TCB can range from 8-32 k bytes and a web page may have 10 or more TCBs.

It is clear that a poorly designed page (that invokes several TCP connections for hyperlinks) can result in a server that very quickly runs out of TCB memory space. It is generally accepted that this is a major cause of the user's perception of poor performance. Reducing the number of objects on a web page and increasing the storage allocated for TCBs in the OS will help the problem.

There is also some ongoing work in the area of a "transactional TCP" protocol extension that will allow one TCP connection to stay "open" for several HTTP transfers, thus minimizing the excessive need for TCB blocks.

When is 100 Mbps not 100 Mbps?

Experiments have demonstrated that WEB transactions over a 100 Mbps link may actually run at rates near ***469 kbps (or less than 0.49 per cent of the link speed***). The experiment used an ATM switch and 2 workstations connected over fiber optic links connected via 100 Mbps pipes. The application was an FTP transfer done using TCP IP whose IP packets were carried by ATM cells. The problem was quite complex but the poor performance was caused by interaction with the TCP IP protocols SWS (silly window syndrome heuristic) and buffer allocation in the transmitting host. It is hard to believe that some simple adjustments in the software and buffer sizes could radically increase performance, but the researchers at Purdue indicated that by simply changing the buffer size in the transmitter (the server), that the overall throughput could be increased to 89 Mbps (close to 90% maximum). The lesson is very simple: protocols and default parameters for protocol buffer sizes are optimized to work in certain types of networks. Oftentimes, these protocols do not scale well with higher bandwidth pipes. Again, legacy equipment, practices, and technology hobble the adoption of higher performance transmission techniques.

Possible Cures

Again, the solution to the overall performance problems of the web are achievable, but at a cost. It is not clear that a business model can be built (by the current transport players) that will significantly improve internet performance at the current cost points (i.e. flat rate pricing, always connected, and no incremental billing per minute).

Loop Solutions

In the local loop, the major high speed options are ISDN, ADSL, and cable modems. Each of the current suppliers of transmission services is saddled with existing infrastructure equipment, technology expertise, business models, regulatory issues, marketing prowess, service reputation, etc. These factors join together to create some rather interesting combinations of partners, press releases, and overall craziness.

The " Cable Guys"

The MSOs are cash strapped and major renovations are required to support bi-directional, high speed access to the internet. The MSOs will also be required to build caching servers, billing systems, maintenance systems, etc., for a product that may deliver only slight increases in performance to the end user (See figure 7). The MSO may have solved the problem of high speed transport only to move the bottleneck further back into the core of the internet. Currently, estimates place the cost of infrastructure upgrades for the MSOs somewhere between 500-1500 dollars per subscriber. This is a tough economic model to live with. Also, the MSOs are plagued with reputation problems, incompatible cable modem standards, etc., that impact the overall deployment, acceptance, and economies of scale for the type of equipment necessary to make this service a reality.

ISDN - the phone company's entry

ISDN provides the 2B 128 kbps transport that surfers crave; in the current format, the B-channels remain connected to the CO switch and consume time slots in a fashion identical to the dial up 28.8 modems (see glossary for discussion on ISDN and time slots). This exacerbates the class V switch blocking problem. Work is underway to produce a standard that allows the ISDN link to go "up and down" depending on the requirements for additional bandwidth. The D-channel (16 kbps) is always "up" in this new standard, thus allowing immediate notification (and delivery) of EMAILs while also allowing the user to 'signal' the net that more bandwidth is needed. Assuming this standard is adopted and deployed, it should allow ISDN to coexist peacefully with the infrastructure for some time. Critics claim that the "always on" nature of the D-channel will move the bottleneck problem to other areas in the switch and outside to the X.25 infrastructure. All of this is speculation and remains to be seen, but clearly, the infrastructure equipment suppliers missed the mark on designing for long holding times, QOS, non-blocking, and low cost.

ADSL (Asymmetric Digital Subscriber Line)

ADSL is a high speed transport technology that uses the existing copper plant around the world (over 700 million trunks are potential targets for ADSL service). The technology uses advanced DSP techniques to recover up to 6 Mbps of downstream data from a copper loop that may reach distances up to 9 kFt. At the central office, the upstream data (operating at 640 kbps) and the downstream data is redirected around the switch using some clever filtering techniques. Therefore, the high speed downstream and low speed upstream information is never connected to the central office switch, thus eliminating class V switch blocking associated with the ADSL transport. Voice calls proceed normally, operating **simultaneously with ADSL data transfer over the same copper wire, and pass through the standard CO switch (see glossary for discussion).**

The redirected data are attached to a device referred to as a DSLAM, digital subscriber line access multiplexer. This device modulates and demodulates the received data from the ADSL line and converts the data to a digital representation suitable for transport in the wide area by the traditional carriers. The unfortunate problem with ADSL technology today is the cost point. Some vendors are offering initial equipment as high as 500 dollars per end point. Again, this is a price point that is far too high to make most RBOC ISP service models profitable. ADSL is also plagued by competing physical layer modulation standards and issues concerning interoperability with other transmission schemes, including ATM. Nevertheless, ADSL is rapidly gaining favor among RBOCs and PTTs as it levers the existing infrastructure (over 600 billion dollars) and the technology curve for ADSL should track the microprocessor / memory system curves for price/performance (as opposed to the cost of backhoes and street repairs required by alternative techniques).

DLCs and AIN

The internet "boom" and the attendant problems have not gone unnoticed by telecommunications equipment vendors. These suppliers are attempting to capitalize on the problems created by the huge and rapid expansion of the internet. For example, DLC providers are offering "smart" loop carrier systems that use AIN (advanced intelligent networks) to 'intercept' the incoming call's destination (AOL, internet, Compuserve, etc.). Once the switch identifies the number, it recognizes that the call is destined for some data service / internet location and assumes that the call will have an excessively long 'holding time.' The call is then routed around the CO switch to some trunking facility that has the capacity necessary to deal with internet traffic. Unfortunately, this solution requires the RBOCs to make a substantial investment in new infrastructure that generates little or no additional revenue.

Wireless

The recent boom in wireless technology has not been wasted on the internet world. The PCS (personal communications system) auctions and explosive growth of cellular technology has refueled the interest in wireless access to data services and the internet. In the ***mobile*** arena, I expect the acceptance of wireless access to the internet to be minimal due to costs and standards issues. However, in the PCS arena, it is possible that aggregation of spectrum from one or more vendors or standard frequency re-use techniques could provide advanced wireless local loop technology (for telephone replacement) as well as internet access at rates that rival ISDN. The economic model is driven by the aggregate of loop bypass and internet access, thus providing a more profitable scenario.

LMDS and MMDS have recently entered the picture as multiple giga Hertz transports for fixed wireless access targeted at the local loop. These technologies certainly have the capability to provide the high speed internet access, but it is not clear that the economic models work for these suppliers. This will not really be known until the spectrum auctions are completed and the total cost of the networks can be computed.

In the Future

It is clear that the demands on the internet will continue to grow for some time. It also clear that the new and more advanced services will begin to appear that demand varying QOS requirements. TCP protocol extensions combined with ATM may sufficiently bolster the infrastructure to continue fueling the growth of the internet. Recent work by Cisco and Ipsilon have highlighted the interaction between switched and routed traffic using ATM transport technology. The detection of "flows" in the TCP data connections allows the routing of time sensitive traffic to networks suited to QOS and minimal hop routes, thus improving overall performance of the network. Some critics say that ATM is a poor transport for hyperlink information as there is little spatial correlation between the hyperlinks activated by a web page access. It is likely that much will be written and researched about RTP and RSVP over the next several months, however, the infrastructure will require modifications to support RSVP (routers and billing, etc.), thus delaying the wide spread acceptance of this technology in the near future.

It is likely that we will continue to see special "deals" cut by various entities in the transport, content, and delivery segments on the internet world. Internet telephony suppliers have finalized deals with IXCs to carry minimal latency voice traffic using TCP-IP and compression to deliver reasonable voice quality and fax service for a fraction of the cost charged by the old guard telephony companies thus continuing the erosion of the artificially high prices charged by IXCs. Likewise, Cisco, Intel and MCI have formed an alliance to deliver multimedia service from end to end using MCI transport, Cisco routing and switching, Cisco local access technology (ADSL looks promising), and Intel computing, including the necessary software support for multimedia.

Conclusions

The internet has been the communications news event of the decade. The rapid growth of the net has strained the overall communications infrastructure from local loops to wide area transport to server and protocol design. It is clear that the entire internet infrastructure must be improved if the astronomical growth rates are to continue in the future. The problem is massive as it involves hundreds of players, millions of users, and billions of dollars of infrastructure controlled by a dozen or so conglomerates. The opportunity to change the landscape of the communications infrastructure is at hand; the rate of change will pale next to the PC industry's meteoric rise, but then there is much that must change within the communications infrastructure that is governed and controlled by many forces. There is opportunity in this new frontier for the patient and the savvy; the ones that can see through the hype and understand the businesses will ultimately prevail; AMD will be a leader in the next generation of this computing paradigm - the fifth generation.

Glossary

ADSL - Asymmetrical Digital Subscriber Loop - a technique of using the existing copper loops to transmit much higher data rates, up to 6 Mbps downstream and 640 kbps upstream by using DSP (digital signal processing) techniques to "squeeze" all available bandwidth out of the copper. ADSL uses an "out of band" modulation scheme that allows telephone calls to occur simultaneously with the ADSL data calls. The ADSL data streams are removed after they enter the CO but before they reach the switch. Typically, the ADSL calls are routed around the switch to a DSLAM (digital subscriber line access multiplexer) that connects to the wide area transport technology of choice.

AIN - Advanced Intelligent Network - a trend towards moving intelligence into the network to provide advanced services and higher reliability. The AIN technology is used to detect internet traffic by recognizing the outbound number that was dialed by a surfer. This call would be routed around any facility that is in danger of being overloaded or "blocked."

CO - Central Office - the location in the telephone network where the local copper loops are concentrated, switched, and powered. Also, connections to the wide area are also provided in the CO.

DARPA - Defense Advanced Research Projects Agency - responsible for the original idea of an "internet."

DLC - Digital Loop Carrier - a technique of locating a switching / concentration center for copper wiring closer to the customer's location. This is particularly useful in remote areas as the copper loops can be installed to serve a cluster of homes or businesses and the aggregated traffic can be "backhauled" to the central office via high speed fiber technology.

DS3 - A high speed communications pipe that operates at 45 Mbps.

GUI - Graphical User Interface - the proverbial "point and click" interface prevalent in MAC and Windows environments.

HTTP - hypertext transfer protocol - the protocol that runs on top of TCP IP used for the transfer of web site information.

ISDN - Integrated Service Digital Network - a technology that was developed in the early 80s to provide digital connectivity from end-to-end across the wide area network. The most prevalent instance of ISDN appears in the local loop as a 2B+D service that offers 2-64 kbps B-channels in conjunction with a 16 kbps D-channel for signaling and low speed packet data. The technology has been rather slow to catch on for a variety of interoperability, tariffing, and infrastructure issues.

ISDN provides higher speed access to the internet but it normally connects to a central office switch not unlike the ones used by dial up MODEM users. Since ISDN consumes 2-64 kbps time slots in the switch, it may actually worsen the switch blocking problem caused by the dial up user (dial-up users only use one 64 kbps time slot). As all switches have a finite number of slots that are typically one tenth the number of total subscribers on the system, the use of ISDN has the potential of being a worse offender for blocking than the dial up user. Work is underway to develop a standard for releasing the B-channel resources when they are not needed. This releases the time slot and makes it available for other users. This is possible in ISDN as the D-channel can set up and release resources in hundreds of milliseconds; this is not possible in the dial up environment.

ISP - Internet Service Provider - A "pay for service" company that offers connections to the internet, EMAIL services, etc., to a surfer. The connection to the ISP is typically provided through the dial up telephone network.

IXC - Inter exchange Carrier including ATT, MCI, Sprint, etc. typically provide transport of voice and data from a central office to another central office.

MSO Multiple systems operator - a cable operator that owns several smaller operating companies that deliver cable service.

NAPs - Network Access Points

OS - Operating System - used by computer to manage system resources and provide services for applications programs.

QOS - Quality of Service - a parameter that is used to define the characteristics expected by an application using communications services. Typically these parameters include bandwidth, jitter, bit error rate, acceptable packet loss rate, etc. The analogy often used by industry pundits is the postal service. First class stamps provide faster service; Federal Express and Express Mail provide guaranteed service at higher cost points. This trend is also gaining favor in the communications arena to assist the transport companies in the handling and routing of various types of information streams from EMAILs to video conferencing.

SOHO - Small Office Home Office

TCP IP - Transmission Control Protocol / Internet Protocol - the defacto routing and transport protocol for internet traffic.

T1 - A medium speed wide area communications pipe that operates at 1.544 Mbps.

WWW - World Wide Web

Net Math

Every HTTP retrieval requires one round trip delay across the network plus TCP set up time of 3 round trip delays. Couple this with a TCP connection termination of 2 round trip delays. Some overlap is possible, but the smallest overall achievable delay is four round trips. Assuming infinitely fast link, and a 300 ms round trip delay, the time required for the page is 1.2 seconds, well under the recommended 3 second “brain fade” guideline discussed earlier, but surprisingly slow for an infinitely fast link.

Interesting Statistics

Monitoring a Cisco 7505 Router between St. Louis and Kansas City, Network Computing reported 58 million flows in 39 hours. Over 55% of these flows were from HTTP. The average flow was 41 seconds long and 68 per cent was idle time.

The average web flow was 13 packets with 3563 bytes; 750 bytes were overhead. The packet flow was insufficient for Van Jacobson “slow start” to have any real impact.

Unfortunately, there are congestion issues on the net that require some tweaks to TCP IP. TCP congestion control reduces throughput until TCP window is fully opened (with no regard for link speed). Experimental results have shown that using TCP to transfer 2 kbyte packets with a 70 ms RT delay results in a throughput of 10 percent of maximum data rate. This number will increase to 50% when the packet size is raised to 20 kbytes. As most of the traffic for web access tends to fall in this category, the predominant flow control mechanisms in TCP are not optimal for the majority of the traffic on the net.

References

Schulzinne, Henning. “World Wide Web: Whence Whither, What Next?” IEEE Network, March, April 1996, 10-17

C. Cuhna, A. Bestavros, and M. Crovella, “Characteristics of WWW Client Based Traces,” Tech. Pub Bu-CS-95-010, Boston University, Computer Science Department, July 1995

Sevcik, Peter. “Designing a high-performance web site,” Business Communications Review, March 1996, 27-31

D. Comer, J. Lin. “TCP Buffering and Performance over an ATM network,” Journal of Internetworking: Research and Experience, vol. 6(1), March 1995 pp1-13

Moskowitz, Robert G. “Why in the world is the web so slow?” Network Computing March 15, 1996, 22-24

Stevens, Richard. TCP IP Illustrated, Volume 1 New York: Addison Wesley Publishing, 1994

Mr. Bell’s Biography

Russ Bell currently serves as the director of communications technology, for the advanced development lab within Advanced Micro Devices, Inc., a fortune 500, California-based semiconductor manufacturer. Mr. Bell is chartered with the responsibility of determining the future direction of AMD's silicon-based products for emerging communications markets. Mr. Bell has also served in a number of capacities for AMD ranging from director, Corporate Strategic Marketing, Communications Technology to the director of North American Field applications engineering during which time he managed a group of 50 plus professional engineering / technical personnel responsible for AMD “design wins” in the PC and communications markets.

Prior to joining AMD in 1984, Mr. Bell served as an Assistant Professor of Electrical Engineering Technology at the Southern Technical Institute in Atlanta, Georgia. Mr. Bell has performed design and consulting work for Scientific Atlanta, Cox Cable, Lockheed-Georgia, and The Department of Defense. Mr. Bell holds degrees in Electrical Engineering Technology from The Southern Technical Institute and a Masters in Information and Computer Science from The Georgia Institute of Technology. Mr. Bell is a member of IEEE and a registered Professional Engineer in the State of North Carolina.

Mr. Bell can be reached via e-mail at: russ.bell@amd.com